

### Claims

What is claimed is:

1. A method for measuring the performance of at least a portion of a computer network, the method comprising the steps of:
  - providing one or more sampling lists, each sampling list identifying one or more nodes representative of a group of nodes of interest;
  - sending signals to a relatively large number of nodes identified on the one or more sampling lists in a relatively short time;
  - receiving signals from responding nodes on the one or more sampling lists;
  - recording raw data indicative of the corresponding sent and received signals;
  - reducing the raw data to data streams indicative of the performance of the nodes; and
  - presenting information derived from the data streams, the information being indicative of the performance of one or more of the groups of interest.
2. The method of claim 1 wherein the providing step further comprises the steps of,
  - gathering data regarding topology of the network,
  - classifying nodes in an organized way based on the gathered data to create one or more lists of nodes for each group of interest,
  - selecting nodes representative of each group of interest, and
  - creating the sampling list for each group of interest from the selected nodes.
3. The method of claim 2 wherein the gathering step further comprises the steps of,
  - obtaining a population survey that lists substantially all of the nodes on the network,
  - extracting a list of main domains from the population survey,
  - recording an example node for each main domain, and
  - identifying the routes through which information travels between the example nodes and a beacon.
4. The method of claim 2 wherein the selecting step further comprises the step of trimming the one or more lists of nodes for the one or more groups of interest by

excluding from consideration those nodes which are not important to the performance of other nodes.

5. The method of claim 4 wherein the selecting step further comprises the steps of, testing the remaining nodes on the one or more lists of nodes for each group of interest to determine reachability of each node, and removing unreachable nodes from the list.
6. The method of claim 1 further comprising the steps of: weeding the one or more sampling lists by removing nodes which are unreachable for a predetermined amount of time, and adding new destinations to the sampling lists as the group of nodes of interest represented by each list increases in size.
7. The method of claim 1 further comprising the step of performing mathematical operations on the data streams to determine performance characteristics representative of a group of nodes of interest for a period of time.
8. The method of claim 1 further comprising the steps of: establishing a threshold value for a change in a performance metric as measured over time; analyzing the data streams using the threshold value to detect any significant changes in the value of a performance metric over time; determining that an event has occurred when a significant change is detected.
9. The method of claim 8 wherein the presenting step further comprises the step of automatically sending an email to interested users whenever an event is detected.
10. The method of claim 1 wherein the presenting step further comprises the step of creating a ratings table showing information indicative of the performance of one or more interesting groups of nodes.
11. The method of claim 10 wherein the information indicative of performance comprises two or more performance metrics.
12. The method of claim 11 wherein the ratings table is user sortable by performance metric.
13. The method of claim 1 wherein the presenting step further comprises the step of creating a graph showing the performance of one or more interesting groups of nodes

over time.

14. The method of claim 13 wherein the step of creating the graph further comprises the step of drawing a shaded region representing a detected performance event.

15. The method of claim 14 further comprising the step of placing tags on borders of the shaded region, the tags comprising information indicative of the nature of the event, and wherein the shaded region further depicts the duration of the event.

16. The method of claim 1 further comprising the step of correcting the data within a data stream by adjusting the values of a performance metric based on geographic distance over which the sent and received signals traveled.

17. The method of claim 1 further comprising the step of correcting the data within a data stream by adjusting the values of a performance metric based on a previously observed metric value for the path over which the sent and received signals traveled.

18. A method for measuring the performance of a computer network, the method comprising the steps of:

providing a survey of network nodes;

creating at least one sampling list representative of at least a portion of the nodes;

performing a process comprising the steps of,

running a sending procedure which sends signals from a beacon to a plurality of nodes on the sampling list, the signals requesting responses from the plurality of nodes,

simultaneously running a receiving procedure which receives return signals from responding nodes, such that the signal being received at any particular time could correspond with any of the signals sent to one of the plurality of nodes, and

simultaneously running a recording procedure which records data characterizing the corresponding sent and received signals;

performing data reduction on the data to create data streams comprising information indicative of the performance of at least a portion of the network; and

presenting information derived from the data streams to a user such that the user is able to draw conclusions about the performance of at least a portion of the network.

19. The method of claim 18 further comprising the steps of,  
determining a size for a chunk that the beacon is capable of running through the process utilizing fast memory exclusively,  
reading a chunk into fast memory before running the sending, receiving, and recording procedures;  
writing the recorded data to slow memory after the sending, receiving, and recording procedures have been completed for the chunk, and  
repeating the reading step, sending, receiving, and recording procedures, and writing step on sequential chunks until all of the chunks have been run through the process.
20. The method of claim 19 wherein the chunk comprises an array of records, each record representing a node to be measured, and wherein the sending procedure further comprises the steps of,  
selecting a record from the chunk,  
transmitting a packet to the destination represented by the record, the packet being encoded with information corresponding to the selected record,  
identifying the record as having a packet in flight, and  
repeating the selecting, transmitting, and identifying steps as long as less than a predetermined number of packets are in flight.
21. The method of claim 20 wherein the sending procedure further comprises the step of incrementing a flying counter any time a packet is transmitted.
22. The method of claim 21 wherein the receiving procedure further comprises the steps of,  
accepting a packet,  
determining whether the packet corresponds with a record identified as having a packet in flight,  
decrementing the flying counter if the packet corresponds to a record identified as having a packet in flight, and  
identifying any corresponding record as not having a packet in flight.
23. The method of claim 20 wherein the receiving step further comprises the steps of,

accepting a packet, and  
determining whether the packet corresponds with a record identified as having a packet in flight, and

identifying any corresponding record as not having a packet in flight.

24. A method of creating a sampling list for measuring performance of a group of nodes of interest on a large computer network, the method comprising the steps of:

obtaining a population survey of the nodes on the network;

gathering data regarding topology of the network;

creating a list of nodes in the group of nodes of interest based on the gathered data;

trimming the list of nodes by excluding from the list those nodes which are not important to the performance of other nodes;

testing the remaining nodes on the list of nodes to determine reachability of each node;

removing unreachable nodes from the list;

categorizing the nodes on the list according to the service type the node provides;

excluding nodes on the list that fall within over-represented service types for the particular emphasis of the group of nodes of interest.

25. A method of correcting data collected about one or more nodes remotely located from a data collection point, the method comprising the steps of:

providing data comprising performance metrics, the data having been collected by a beacon at the collection point by sending signals to and receiving signals from one or more nodes;

determining a great circle distance over which the signals traveled between the collection point and one or more nodes;

calculating the time required for the signal to traverse the distance using the speed of light within the link and the determined distance for one or more nodes to find a geographic distance factor; and

subtracting the appropriate geographic distance factor from the metric values within the data to provide corrected data.

26. A method of correcting new data collected about one or more nodes remotely located from a beacon, the method comprising the steps of:

providing pre-collected data, the pre-collected data having been collected by the beacon by sending signals to and receiving signals from one or more nodes, and the data comprising performance metrics for a predetermined length of time;

determining a baseline metric value for the pre-collected data to find a correction factor;

subtracting the correction factor from the metric values within the new data to provide corrected data.

27. A general method for evaluating performance data of a computer network system, comprising:

receiving performance data measured from a computer network system;

selecting an appropriate event detection method to evaluate the performance data, the selected event detection method adapted to compare the performance data to threshold parameters to determine whether an event has occurred;

sending a signal indicating an event has occurred whenever an event is detected by the appropriate event detection method.

28. The method for evaluating performance data of a computer network system according to claim 27, wherein the selected event detection method comprises:

measuring the average value of performance data in a window representing a predefined number of data points;

comparing the average value of the data points within the window to determine whether it falls within a threshold range; and

recording an event if the average value of the performance data falls outside of the threshold range.

29. The method for evaluating performance data of a computer network system according to claim 27, wherein the selected event detection method comprises:

selecting a ratio threshold, an absolute minimum threshold value, and an absolute minimum difference threshold;

determining the number of consecutive data points in a sequence making up a sliding window;

selecting a first series of data points as a first sliding window;  
calculating a first average value of the data points in the first sliding window;  
selecting a second series of data points corresponding to the number of  
consecutive data points comprising a sliding window as a second sliding window,  
wherein the second sliding window comprises the data points in the first sliding  
window minus the oldest data point in the first sliding window, plus a new added data  
point that is the next data point in a sequence;

calculating a second average value of the data points in the second sliding  
window; and

recording an event if the ratio of the second average to the first average is  
greater than or equal to the ratio threshold, the second average exceeds an absolute  
minimum threshold value, and the difference between the average value of the second  
window and the average value of the first exceeds an absolute minimum difference  
threshold.

30. The method for evaluating performance data of a computer network system  
according to claim 27, wherein the selected event detection method comprises:

setting a range of line slope values as a threshold range for event detection;  
setting an array of weights in a filter kernel to be applied to performance data  
points;

adding data points to a filter window until a predetermined number of data  
points has been added to the filter window;

applying the filter kernel to the data points in the filter window to produce  
filtered data points;

adding filtered data points to a line fit window until a predetermined number of  
filtered data points is added to the line fit window;

fitting the best least squares line to the filtered data points in the line fit  
window;

determining the slope of the fitted line;  
comparing the slope of the fitted line to the threshold range of line slopes; and  
recording an event if the slope of the fitted line falls outside of the threshold  
range.

31. The method for evaluating performance data of a computer network system according to claim 27, wherein the selected event detection method comprises:

- setting a range of line slope values as a start slope threshold for detecting the start of an event;

- setting a range of line slope values as an end slope threshold for detecting the end of an event;

- setting an array of weights in a start filter kernel to be applied to performance data for detecting the start of an event;

- setting an array of weights in an end filter kernel to be applied to performance data for detecting the end of an event;

- setting the initial state of the system to "not\_event";

- adding data points to a start filter window until a predetermined number of data points has been added to the start filter window whenever the system state is "not\_event";

- producing start filtered data points by applying the start filter kernel to the start filter window whenever the system state is "not\_event";

- fitting a best least squares line to the start filtered data points whenever the system state is "not\_event";

- determining the slope of the line fitted to start filtered data whenever the system state is "not\_event";

- comparing the slope of the fitted line to the start slope threshold whenever the system state is "not\_event";

- recording the start of an event if the slope of the fitted line falls outside of the threshold range for detecting the start of an event whenever the system state is "not\_event";

- changing the state of the system to "event" if the start of an event is recorded;

- adding data points to an end filter window until a predetermined number of data points has been added to the end filter window whenever the system state is "event";

- producing end filtered data points by applying the end filter kernel to the end filter window whenever the system state is "event";



fitting the best least squares line to the end filtered data points whenever the system state is "event";

determining the slope of the line fitted to end filtered data whenever the system state is "event";

comparing the slope of the fitted line to the end slope threshold whenever the system state is "event";

recording the end of an event if the slope of the fitted line falls within the threshold range for detecting the end of an event and if the variance of the data points within the end filter window falls within a given threshold range whenever the system state is "event"; and

changing the state of the system to "not\_event" if the end of an event is recorded.

32. A method of combining data streams comprising:

providing two or more input data streams, the data streams representing data collected over a time interval, and each data stream comprising one or more records, each record including a start time indicating the time at which a collection process began collecting the data within the record and a value for one or more metrics;

defining sequential time slices, each time slice comprising a piece of the interval and being the duration of the interval divided by a frequency;

sorting the records from the two or more data streams by the start time;

allocating each record into the time slice in which the start time for each record falls;

calculating one or more statistics for each time slice, each statistic representing the value of one or more metrics for the records allocated to each time slice; and

outputting a data stream comprising records which include the calculated statistics.

33. The method of claim 32 wherein the calculated statistics comprise medians.

34. The method of claim 32 further comprising the step of:

defining an offset,

wherein the first time slice begins at the first time frequency after the interval begins plus the offset.

35. The method of claim 32 wherein the formats of the output data stream and each of the input data streams are identical.

36. A method of selectively notifying interested users of a performance event, the method comprising the steps of:

receiving a signal indicating that an event has been detected and the type and magnitude of the event;

composing an email message indicating that an event has been detected and the type and magnitude of the event;

comparing the indications to records within a file comprising conditions in which users desire to be notified of performance events;

automatically sending the email to users whose records included a condition met by the indications.

37. The method of claim 36 wherein the composing step further comprises the step of including in the email an auto-reply option which a user can quickly use to indicate that the user does not want to be notified of similar events in the future, and the method further comprising the step of updating the conditions in the responding user's record in the file.

38. A method of displaying performance characteristics of at least a portion of a network, the method comprising the steps of:

providing values for metrics indicative of the performance of at least a portion of the network; and

producing an image including cells organized in columns and one or more rows, each column representing a performance metric, each row representing at least a portion of the network, each cell including values for a metric as measured over one or more time intervals.

39. The method of claim 38 wherein the order in which portions of the network are listed is user sortable by metric.

40. A method of displaying performance characteristics of at least a portion of a network, the method comprising the steps of:

providing values for metrics indicative of the performance of at least a portion of the network;

producing an image including a time versus metric graph indicative of the performance of at least a portion of the network;

drawing a shaded region representing a detected performance event on the graph.

41. The method of claim 40 further comprising the step of placing tags on borders of the shaded region, the tags comprising information indicative of the nature of the event, and wherein the shaded region further depicts the duration of the event.

42. A method of collecting data from a large number of nodes on a computer network in a short time, the method comprising:

providing a sampling list comprising information indicative of the addresses of nodes about which data is to be collected; and

performing a process comprising the steps of,

running a sending procedure which sends signals from a beacon to a plurality of nodes on the sampling list, the signals requesting responses from the plurality of nodes,

simultaneously running a receiving procedure which receives return signals from responding nodes, such that the signal being received at any particular time could correspond with any of the signals sent to one of the plurality of nodes, and

simultaneously running a recording procedure which records data characterizing the corresponding sent and received signals.

43. The method of claim 42 further comprising the steps of,

determining a size for a chunk that the beacon is capable of running through the process utilizing fast memory exclusively,

reading a chunk into fast memory before running the sending, receiving, and recording procedures;

writing the recorded data to slow memory after the sending, receiving, and recording procedures have been completed for the chunk, and

repeating the reading step, sending, receiving, and recording procedures, and writing step on sequential chunks until all of the chunks have been run through the process.

44. The method of claim 42 wherein the sent signals are pings.

45. The method of claim 44 wherein the number of pings in flight is limited to a predefined value, pings in flight being those pings which have not timed out or for which a response has not yet been received.

46. A method of collecting topological information about a large number of nodes on a computer network in a short time, the method comprising:

providing a sampling list comprising information indicative of the addresses of example nodes; and

performing a process comprising the steps of,

running a sending procedure which sequentially sends a predetermined number of signals from a beacon toward each of a plurality of nodes on the sampling list, the sequentially sent signals being sequentially encoded with a time to live parameter,

simultaneously running a receiving procedure which receives return signals from responding hops along the route to each destination, the hops reporting an expired time to live, such that the signal being received at any particular time could correspond with any of the signals sent toward the plurality of nodes, and

simultaneously running a recording procedure which records data characterizing the corresponding sent and received signals.

47. The method of claim 46 further comprising the steps of,

determining a size for a chunk that the beacon is capable of running through the process utilizing fast memory exclusively;

reading a chunk into fast memory before running the sending, receiving, and recording procedures;

writing the recorded data to slow memory after the sending, receiving, and recording procedures have been completed for the chunk, and

repeating the reading step, sending, receiving, and recording procedures, and writing step on sequential chunks until all of the chunks have been run through the process.